DAILY AND INTRADAILY STOCHASTIC COVARIANCE: VALUE AT RISK ESTIMATES FOR THE FOREIGN EXCHANGE MARKET

George Kormas

A Thesis
In
The Faculty
of
Commerce and Administration

Presented in Partial Fulfilment of the Requirements
for the Degree of Master of Science in Administration at
Concordia University
Montreal, Quebec, Canada

May 1998

0-612-39466-2

ABSTRACT

Daily and Intradaily Stochastic Covariance:
Value at Risk Estimates for the Foreign Exchange Market

George Kormas

The importance of time varying volatility in securities prices (e.g. GARCH) has by now been amply established in the literature, both in terms of the magnitude and pervasiveness of the phenomenon, and in terms of its significance for risk management in institutional portfolios. Less attention has been devoted to multivariate conditional heteroskedasticity, in spite of the fact that securities are typically held in portfolios rather than in isolation. Recently, Kroner and Ng (1995) have introduced a method for nesting the four most commonly used multivariate GARCH models, allowing for comparative tests of the performance of the models. We propose to apply the Kroner and Ng technique to both daily and intradaily returns on foreign exchange rates, to obtain performance estimates. These conditional covariances will then be used to calculate value at risk (VaR) forecasts for foreign currency portfolios. Daily and intradaily VaR forecasts will be evaluated and compared.

DEDICATION

Προς τους γονεις μου,
Ευσταθιος και Κωνσταντινα,
που αφιερωσαν τη ςωη τους στα παιδια τους.

Table of Contents

1. Introduction 1

2. Literature Review 5
- **2.1.** Value at Risk 5
  - **2.1.1.** Historical Simulation 7
  - **2.1.2.** Structured Monte Carlo Simulation 9
  - **2.1.3.** Variance-Covariance Method 10
  - **2.1.4.** Value at Risk – Comparative Studies 12
- **2.2.** Conditional Heteroskedasticity in the Foreign Exchange Market 14
  - **2.2.1.** Early Research 15
  - **2.2.2.** The Multivariate Models 18
  - **2.2.3.** Intradaily Foreign Exchange Volatility 21

3. Empirical Results 25
- **3.1.** Description of the Study 25
- **3.2.** The Dataset 26
- **3.3.** Preliminary Statistics 27
- **3.4.** The General Dynamic Covariance Model 29
- **3.5.** Model Estimation Results 32
  - **3.5.1.** General Dynamic Covariance Model – Results 32
  - **3.5.2.** Constant Correlation Model – Results 34
- **3.6.** Comparison of Models 36

4. Concluding Remarks 40

5. References 42

6. Appendices 46

1. INTRODUCTION:

Much research has been devoted to the notion of market volatility since Harry Markowitz's pioneering work on the risk-return relationship in the early 1950's. Since this relationship was first introduced to the financial community, academics and practitioners alike have addressed the ever-present role that volatility plays in financial management. Furthermore, the globalization of the world economies as well as the sophistication of the financial contracts now in existence has necessitated a better understanding of market volatility for effective portfolio management.

In recent years, heightened awareness of market volatility and market risk has been brought about by some highly publicized losses, Orange County and Daiwa to name a few. On account of these and others, financial institutions have placed greater importance on their risk management practices. Better management of financial risks that are affected by market volatility requires better quantification of volatility itself. Therefore, the risk management community has begun to scrutinize its risk quantification models and can no longer take for granted the assumptions underlying volatility measurement.

Particular attention to risk quantification models has been paid by the regulators of the financial community, for it is the regulatory bodies that must ensure the financial stability of the institutions to which investor wealth is entrusted. An example of regulatory initiatives in risk management is portrayed by the recent legislation imposed by federal agencies on the chartered banks, as proposed by the Basle Committee on Banking Supervision (1997). The Basle Committee's document has as its main purpose to outline the principles by which the banking community must operate and manage its affairs. One of the focal points of this report is on the management and accurate measurement of portfolio market risk. Specifically, the accord stipulates that regulatory authorities within each country develop a means of assessing and validating the internal models used by the chartered banks in their risk management practices. To accomplish this, a standardized risk measure was needed to set risk based capital requirements across different institutions. It was this need that gave rise to the tool now referred to as Value at Risk.

Value at Risk, or VaR, is a model adopted by the risk management community as a means of estimating maximal portfolio devaluation associated with adverse market movements. Not only has VaR been accepted as an excellent risk-monitoring tool, but it is used extensively in capital allocation decisions as well. One of its appealing features is that it is a summary measure of portfolio, or firm-wide, market risk. It expresses in a single number the amount of money that a firm may lose during a given time horizon due to adverse market movements.

In order for a VaR model to estimate the "capital-at-risk" of a portfolio, certain assumptions pertaining to market movements must be made. Specifically, assumptions must be made on how the assets, or market factors, that constitute the portfolio evolve over time. Clearly, any model used to describe asset or market factor returns requires an assumption, either explicit or implicit, on the volatility of these variables. In the case of VaR, as in other areas of finance such as asset pricing or hedging, the choice of volatility assumption is crucial. The dependency of VaR estimates on the underlying volatility assumption has been extensively documented in the risk management literature.

Much of the research on market volatility has been focused on its temporal dynamics. Since the introduction of the univariate ARCH (Engle 1982) and GARCH (Bollerslev 1986) processes, models of time-varying volatility have gained increased attention in the literature. Unfortunately, much less attention has been paid to time-varying covariance effects, regardless of the fact that securities are typically held in portfolios rather than in isolation. From a risk management perspective of an institutional portfolio, the effects of covariance on portfolio volatility play an important role for the reason stated above. Therefore, in order for a risk management practice to be effective in its mandate, the proper measurement/estimation of the entire covariance matrix must be addressed.

From a purely academic standpoint, the estimation of the covariance matrix for a portfolio of assets presents an interesting study. A consistent and robust estimate would provide some empirical evidence on the behavior of not only the variability of individual market factors, but also on the co-movement of these factors and therefore would contribute to our understanding of the price process of the portfolio. From a practitioner's perspective, the value associated with such a study is quite different. The use of a specific volatility model in the context of VaR is constrained by the model's ability to provide accurate and reliable volatility forecasts. In other words, the predictive power of a particular model is of utmost importance. In what follows, we will attempt to address both of these issues in the following manner: First, we will study an area that has received little attention in the literature, namely, heteroskedasticity in covariance. To accomplish this task, we will be using a multivariate parameterization of the covariance matrix recently proposed by Kroner and Ng (1995) entitled the General Dynamic Covariance (GDC) model. This model is a hybrid of the four most commonly used multivariate GARCH models while still allowing for time variation in the covariance terms. In addition, we will contrast the results of the GDC to the less sophisticated Constant Correlation model of Bollerslev (1990). As part of the study, we will derive estimates of the covariance matrix at the daily, as well as at various intradaily levels. The aim of this approach is to take advantage of the high frequency price data that is now readily available and draw some inference on the temporal relationship between low and high frequency estimates of the covariance matrix. Second, we will be using the estimated variance-covariance matrices to determine the VaR for a hypothetical foreign currency portfolio. These estimates will then be compared to results obtained by JP Morgan's RiskMetrics™ methodology, which has developed into the benchmark for VaR estimation among risk management professionals.

The structure of this paper is as follows: the next section is devoted to a review of the literature on both VaR and conditional heteroskedasticity in the foreign currency market. We will outline the most commonly used VaR and multivariate GARCH models as well as summarize recent empirical results in studies of the market microstructure. The third section will be used to describe the study in detail, as well as provide the empirical results of the estimations and the data analysis. Furthermore, we will compare the predictive performance of the two volatility models as well as contrast the VaR estimates based on the GDC and RiskMetrics™ models. Section four is reserved for a summary discussion of the paper and some concluding remarks.

2. LITERATURE REVIEW

2.1. Value at Risk

In spite of the fact that VaR is a relatively new topic in risk management, much research has already been devoted to exploring its usefulness and shortcomings. Although a myriad of VaR models do exist today, each with differing assumptions and methodologies, they all attempt to address the same issue – what is a portfolio's expected decline in value with a pre-described probability. More formally, VaR may be considered a summary measure that defines the expected maximum loss over a target horizon within a given confidence interval (Jorion 1997).

Although a change in the value of a portfolio can be attributed to numerous different risk components, VaR typically attempts to estimate portfolio devaluation from the standpoint of market risk. Market risk involves the uncertainty of future earnings resulting from changing market conditions (e.g. prices, rates) (JP Morgan 1995). All of the VaR models used today to monitor the market risk of a portfolio can be classified in one of two categories, they are either 1) Analytical models or 2) Simulation models.

The analytical models used in VaR estimation are functions of the position sensitivity of the portfolio to the underlying market factors and to the estimated changes in these market factors. A good interpretation of the latter is given in Zangari (1997),

estimated value change = f(position sensitivity, estimated rate/price change). (1)

Essentially, VaR estimates that rely on analytical approaches amount to local approximations of the change in portfolio value in much the same way as a Taylor series expansion is used to approximate the change in function value given a change in the independent variables. Furthermore, the degree of sensitivity used in any analytically based VaR model will vary, ultimately depending on the composition of the portfolio itself. For the case of strictly linear payoffs[1], the first order moment is adequate (delta). In the presence of bonds or derivative securities with non-linear payoff structures, the second order moment is usually included (gamma). Using the above approximation results, one then can produce an estimate of portfolio value change given changes in the underlying market factors.

Simulation methods, more commonly referred to as full valuation methods, rely on revaluing the portfolio under different market conditions or scenarios. Specification of these scenarios varies by VaR model and may include both subjective and objective information with regards to market movements. In spite of the fact that different simulation models use different approaches in generating these scenarios, the end result remains the same: to produce an estimate of the change in portfolio value given the sensitivity of the portfolio to the market factors that affect it.

From the above description of analytical and simulation models, it is evident where they differ. It is in their measurement of the market risk of a portfolio of assets. The analytical approach uses well defined functional characteristics of the pricing formulae to approximate the effect of market movements, whereas the simulation-based VaR models use the computationally more intensive approach of re-calculating the value of the portfolio under specified market conditions. In fact, the aforementioned classes of VaR models do differ in another respect. Specifically, they differ in their treatment of the estimated changes in the underlying market factors that affect portfolio value by imposing different assumptions on the volatility of these market factors. This particular difference in

---

[1] E.g. equity or foreign exchange positions with no optionality.

VaR models has been the element that has received the most attention in the risk management literature. Furthermore, it has been the driving force behind much of the research that has been conducted in the search for a better method of estimating portfolio or firm-wide VaR. What follows is a description of the three most widely used VaR models in risk management today. The first two fall under the category of simulation models, whereas the last one is an analytical model and is the main focus of this paper.

2.1.1. Historical Simulation

Historical simulation is by far the most straightforward of all the VaR models to implement and to understand (Linsmeier 1996). As its name implies, this method uses the full valuation approach to VaR estimation, but its relative ease of use stems from the way in which changes in the underlying market factors are estimated. Historical simulation methods make no explicit assumptions as to the probabilistic nature of asset returns or on the market factors influencing prices, but instead, rely solely on the recorded historical movements of the markets (Zangari 1997). Clearly, the only assumption that proponents of the Historical simulation method make, is that the past is a good guide for future price movements.

In order to calculate the VaR for a portfolio of assets, choices must be made with respect to the holding period and confidence level (or critical region) for which the resulting VaR is representative. When using historical simulation, an additional choice must be made concerning the historical period from which estimated changes in the underlying factors are derived. Essentially, implementation of this method subjects the portfolio to historical changes in market rates and prices over the pre-specified period. From this, we are able to build an empirical distribution of portfolio gains and losses for the next investment horizon retrospectively (RiskMetrics™ 1995). The resulting distribution can then be used to read off the loss that is exceeded x% of the time, where x is the chosen size of the critical region. This loss would then be the portfolio VaR for the next time period.

The methodology outlined above is the most basic of all the full valuation approaches. Its strengths are in its simplicity, in the objective manner with which the distribution is generated and in its independence of statistical and/or distributional assumptions. Its disadvantage lies in the single sample path approach it takes to VaR estimation. The underlying assumption is that the factors affecting portfolio value will change in exactly the same fashion in the next period as they did in the previous one. Clearly, this limitation precludes the use of sensitivity analysis to ascertain how the VaR would change under different assumptions of market factor volatility (Simons 1996).

The restrictive nature of the historical simulation model described above has led practitioners to a different, but related, full valuation model called stress testing. This approach requires the specification of numerous sets of market scenarios upon which the valuation of the portfolio is desired. Using each of these sets, a range of plausible portfolio returns can be generated. Assigning an occurrence probability to each return fully specifies an empirical distribution of portfolio value, from which, the VaR can be obtained as in the historical simulation approach (Jorion 1997).

The stress testing method, although alleviating the problem inherent in historical simulation, suffers from its own drawbacks. First, the changes in the market factors are determined in a manner that is completely subjective since the scenarios are not necessarily selected to be representative of current or future market conditions. Furthermore, There is no guideline imposed by the method as to the inclusion or exclusion of one scenario over another. Second, even if all the improbable market outcomes could be removed from the analysis leaving only the most probable ones, the act of arbitrarily attaching a probability to each return in the remaining set would most certainly bias the estimate of portfolio VaR.

2.1.2. Structured Monte Carlo Simulation

It is clear how the drawbacks of both the historical simulation and the stress testing methods could lead to misrepresentation of portfolio VaR. The need for a more comprehensive VaR methodology led to the development of another full valuation approach called Structured Monte Carlo Simulation, or simply, Monte Carlo Simulation. Monte Carlo Simulation can be viewed as a generalization of the previous two methodologies discussed. By removing the need for user defined scenarios and by introducing more flexibility into the model, the Monte Carlo approach is a more systematic approach to VaR estimation.

In contrast to the other two full valuation methods, Monte Carlo introduces numerous scenarios to the VaR model through the use of explicit statistical assumptions. These assumptions are expressed in the form of probability distributions for changes in the underlying market factors. Once this is done, estimated changes in the market factors can be simulated using standard numerical techniques and the resulting distribution of portfolio profits and losses can be obtained. Repeating this procedure numerous times produces a set of profit and loss distributions, from each of which an estimate of portfolio VaR can be derived. Since the relationship between simulated paths and VaR estimates is one to one, the expected portfolio VaR using this methodology reduces to the arithmetic (possibly weighted) average of the set of VaR estimates obtained from the set of simulated paths. It should be clear to the reader at this point that Monte Carlo Simulation allows for a wider spectrum of possible portfolio profits and losses in the resulting distribution since it accounts for many different portfolio paths rather than just a select few (RiskMetrics™ 1995).

Clearly, Monte Carlo Simulation allows for a more comprehensive approach to VaR estimation. Its flexibility, as well as its ability to capture potential movements in the capital markets, makes it a very appealing alternative to the other static full valuation methods. Unfortunately, the advantages offered by Monte Carlo Simulation are overshadowed by the one disadvantage that deters many financial institutions from implementing it. The level of sophistication required by both the end-user and the necessary information processing technology is quite high. For example, if one were to generate 1,000 sample paths for a portfolio of 1,000 assets, the total number of necessary valuations would come to 1,000,000 (Jorion 1997).

2.1.3. Variance-Covariance Approach

As an alternative to the computationally intensive or inflexible full valuation VaR models, analytical approaches to VaR estimation have been developed. The most widely used analytical method is referred to as the Variance-Covariance approach. As its name implies, this method relies on the existence of the covariance matrix for the portfolio factors and makes use of the concepts of modern portfolio theory (Markowitz 1952) to determine expected changes in portfolio value.

The most discernible difference between the full valuation models already discussed and the Variance-Covariance approach is in the estimation of changes in the underlying factors. In contrast to the other models, the Variance-Covariance method is based on the assumption of a joint distribution[2] between the portfolio factors. Given this distributional assumption and an estimate of the covariance matrix (which is derived from historical time series), we are then able to fully specify the distribution of portfolio profit/loss and determine the percentile corresponding to the portfolio VaR (Zangari 1997).

In order to estimate the portfolio VaR under the Variance-Covariance approach, we must first determine the variance of the portfolio. Referring to the method introduced by Markowitz, the portfolio variance is given as:

$$\sigma_P^2 = \sum_{j=1}^{n} w_j^2 \sigma_j^2 + \sum_{j=1}^{n} \sum_{\substack{k=1 \\ j \neq k}}^{n} w_j w_k \sigma_{jk} \qquad (2)$$

where n denotes the number of factors in the portfolio and w_j is the amount invested in the factor j. Once the portfolio variance is determined, we can make use of the symmetry property of the normal distribution to ascertain the portfolio VaR. Therefore, for a given level of confidence α, the corresponding percentile can be obtained from standard normal tables. Thus producing a VaR estimate of:

$$VaR_P = \sigma_P z_\alpha \qquad (3)$$

Given the symmetry of the normal distribution as well as the fact that sums of normal distributions are also normal, it is clear how efficient a framework the Variance-Covariance method is in estimating portfolio VaR. Many studies however have criticized this approach. The source of the criticism is rooted in the normality assumption of factor returns. It has been extensively documented in the finance literature that factor, or asset

[2] Previous studies usually employ a joint normal distribution.

returns, exhibit much fatter tails (i.e. kurtosis), hence much more peaked, than what is described by the standard normal distribution (Bollerslev 1986).

2.1.4 Value at Risk - Comparative Studies

Since its inception, VaR has been held in high regard as an effective risk management tool. Its ability to summarize risk based capital requirements in a single number has led the risk management community to adopt this approach as the de facto standard in risk measurement (Jorion 1997). The only question that remains unanswered is which method produces better estimates of portfolio VaR. From the description of the VaR models given above, it is clear that substantial differences between models exist in their volatility assumptions. Moreover, perhaps it is this difference that precludes any of the existing studies from providing any definitive results on model superiority. The purpose of this section is to explore a few of the studies that have been undertaken as well outline the conclusions that have been drawn.

As regulators imposed deadlines on the chartered banks for the adoption of VaR risk reporting, many practitioners and academics studied the available methodologies in search of the most comprehensive VaR model. From this search, a host of results have emerged all echoing the same theme: VaR estimates are extremely sensitive to the choice of assumptions underlying the model and that there is no one model that is superior to the others in all circumstances.

In Beder (1995), the author investigates three portfolios and estimates VaR using three methodologies: 1) Historical Simulation, 2) Monte Carlo with a RiskMetrics™ covariance matrix and 3) Monte Carlo with a BIS/Basle covariance matrix. Furthermore, the author varies the holding period upon which the VaR estimates were computed as well as the historical window from which the state variable covariance matrix is derived. The results provided indicate large discrepancies between, as well as within the various methodologies, implying that VaR is highly sensitive to the choice of assumptions. In addition, the discrepancies did not follow a discernible pattern as the complexity of the portfolio was increased from a fixed income portfolio to a portfolio consisting entirely of equity index options.

In a different study, Hendricks (1996) estimated VaR on a hypothetical foreign currency portfolio and found different results. The author contrasts the variance-covariance approach (both equally weighted and exponentially weighted moving average estimates of volatility were used) with the historical simulation approach and reports that, for estimates with varying historical windows, the two approaches do not produce VaR estimates which differ substantially on average. One of the interesting findings of this study is that volatility models that use longer historical windows produce less variable VaR estimates over time, an intuitively appealing result. This is evident since finite-window moving average models produce "shadows" which generate non-uniform estimates of volatility, a phenomenon which is well documented in the literature (Diebold 1987).

In a recent study on non-linear VaR, Fallon (1996) tested the performance of various variance-covariance methods on a derivatives portfolio by incorporating both delta and gamma components, while maintaining a heteroskedastic covariance structure (multivariate GARCH) for the underlying state variables. By including the second order moment in the Taylor's series approximation to the change in portfolio value, Fallon finds the resulting distribution of the mark-to-market profit and loss to be a non-central χ^2

distribution. In a similar paper, Jones and Schaefer (1997) later confirmed this result. In spite of the fact that Fallon employed a restrictive constant correlation model for the state variables, he finds that this multivariate parameterization outperformed the other VaR models tested, even though out-of-sample tests of the volatility assumption seem to provide little or no predictive power[3].

Given the preliminary findings of the first two studies described above, it is not difficult to appreciate the dependency of VaR on the underlying assumptions. Two main points arise from these studies: the first is that different approaches to VaR estimation may produce different results. This is not difficult to envision since the various VaR methodologies explored do differ in their volatility assumptions as well as in the way that they use historical data. Second, these studies show that comparisons of VaR estimates based on similar methodologies may also produce different results. It is important to note however that the volatility estimates employed in these studies were not selected by virtue of their consistency or their efficiency[4], but rather, they were selected in an ad hoc manner (Hendricks 1996). In contrast to the first two studies, Fallon's results indicate that a within model comparison, when coupled with a statistically appropriate measure of volatility, may indeed provide insight as to the relative performance of a particular class of VaR methods.

2.2. Conditional Heteroskedasticity in the Foreign Exchange Market

Since the introduction of ARCH (Engle 1982) and GARCH (Bollerslev 1986), the empirical studies that have emerged to investigate the usefulness of these models have been exhaustive. Mandelbrot (1963) first documented certain empirical regularities in asset

[3] This will be addressed further in Chapter 3.
[4] This is used in the context of statistical efficiency and consistency.

returns and it was those findings, as well as subsequent ones, that motivated the development of the (G)ARCH class of time series models. The excess kurtosis, as well as the temporal persistence of volatility (i.e. volatility clustering) evident in many empirical studies of financial time series raised questions as to the validity of the standard Box-Jenkins parameterization. Because of this, researchers searched for a volatility model that described the return generating process in a less restrictive fashion than the assumption of homoskedasticity allowed for (Bollerslev 1993). One of the areas of the capital markets where the (G)ARCH assumption has been extensively tested has been in the foreign exchange market.

2.2.1. Early Research

The unimodal as well as the leptokurtotic nature of spot foreign exchange returns have been documented as far back as Burt (1977), Westerfield (1997) and Rogalski (1978), but since then, the latter characteristic has received much more attention from the research community. Two possible explanations have been set forth to explain this apparent idiosyncrasy of the foreign exchange market. The first implies that the data are drawn from a fat-tailed stationary distribution whereas the second assumes a temporal dependency in the process generating the returns (Hsieh 1988). The focus of the (G)ARCH class of models have been to test the latter hypothesis.

In his original work, Bollerslev (1986) introduced the GARCH model by building on the results found by Engle (1982). Furthermore, the development of both these models were driven by similar motivations; that given the empirical evidence of heteroskedasticity in the asset returns, traditional homoskedastic time series models were unable to capture the stylized facts of the short run dynamics in the foreign exchange market (Bollerslev et al. 1992). The general form of the GARCH(p,q) model is given below:

$$\varepsilon_t \mid \psi_{t-1} \sim N(0, h_t) \qquad (4)$$

$$h_t = \alpha_0 + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2 + \sum_{k=1}^{p} \beta_k h_{t-k} \qquad (5)$$

In this model, ε_t is a real valued discrete-time stochastic process and ψ_t is the filtration generated by the underlying time series containing all information through time t (Bollerslev 1986). The appealing feature of this model is how the parameterization is well suited to the properties exhibited by financial time series. Clearly, a time series representation with a GARCH variance structure allows for time-variation (i.e. heteroskedasticity) in the return generating process and is a conditional measure on all of the information realized. Moreover, the ARMA structure in the variance equation is meant to capture both the clustering (i.e. the momentum in the conditional variance) that occurs during highly volatile periods as well as the leptokurtosis in the return distribution (Baillie 1989). In other words, GARCH processes were developed to formalize the observed reality that changes, large or small (and of either sign), seem to be followed by further large or small changes respectively. Furthermore, the observed excess kurtosis in spot foreign exchange returns as documented by Westerfield (1977) and others, indicates that a linear GARCH structure is appropriate since the unconditional density produced is fat-tailed whereas the conditional densities remain normal (Diebold 1987). Therefore, the conditional moment structure of GARCH gives us the ability to model the contiguous periods of volatility and stability evident in many markets, including the currency markets.

Although the generalized model allows for lags up to p and q for the variance and the squared innovations respectively, many empirical studies have found that a GARCH(1,1) parameterization is sufficient to describe the heteroskedasticity in foreign exchange returns (Bollerslev et al. 1992). More specifically, Hsieh (1989) found that the standard GARCH(1,1) as well as an EGARCH were extremely successful at removing conditional heteroskedasticity from daily foreign exchange movements. Baillie et al. (1989) also confirmed this result. In addition, Bollerslev et al. (1992) highlight an important empirical finding concerning the presence of ARCH effects, namely that the presence of ARCH tends to weaken with less frequent sampling – with the dissipation of heteroskedasticity becoming clear at weekly levels and almost non-existent at monthly intervals.

In addition to capturing heteroskedasticity in the foreign exchange market, the empirical studies on GARCH processes have documented well the persistence (i.e. momentum) in conditional volatility. Numerous researchers including Lamoureux et al. (1990) and Bollerslev et al. (1992) confirmed this result. Hsieh (1989) found that the persistence in daily foreign exchange rates was very close to unity for many currencies and was statistically indistinguishable from one in many cases. Given these results, the GARCH class of volatility models has enjoyed much success in specifying the volatility of spot foreign exchange returns.

One of the major criticisms of GARCH processes is rooted in its inherent structure. Given its ARMA representation, many studies have documented GARCH's inability to capture sudden structural shifts in volatility (Lamoureux et al. 1990). In other words, GARCH imposes a smooth transition from highly volatile environments to ones with low volatility (and vice versa), even though in reality, transitions such as these do not occur as smoothly. The incongruency between GARCH estimates of volatility and actual market movements has a large implication with respect to VaR models employing a GARCH process for the state variables. Specifically, if GARCH is unable to capture or model sudden shifts in the markets, and VaR is meant to capture the losses associated with these shifts which are low probability events, then are GARCH processes even appropriate in this circumstance? This inconsistency has many researchers investigating different ways of modeling these sudden shifts; one of the more recent methodologies is that of Extreme Value Theory. For a more detailed discussion of EVT, the interested reader is referred to Danielsson et al. (1997) where the authors explore its potential uses in modeling tail events in the context of VaR.

2.2.2. The Multivariate Models

Much research has been devoted to evaluating the usefulness of univariate GARCH time series models. Surprisingly though, very little work has been conducted in a multivariate GARCH setting, regardless of the indisputable need of a statistically sound measure of covariation between assets in many financial applications. Some examples would include portfolio allocation problems as well as derivatives' pricing when there is more than one underlying security (e.g. spread options). Fortunately, many of the techniques used in the univariate GARCH setting can be easily extended to a multivariate framework (Engle et al. 1996). What follows is a summary of the four most commonly used multivariate conditional heteroskedasticity models.

The simplest multivariate GARCH model to understand is the vector GARCH (or VECH) model as proposed by (Bollerslev et al. 1988). The i,j^{th} entry of the covariance matrix under the VECH parameterization is:

$$h_{ij,t} = \omega_{ij} + \beta_{ij} h_{ij,t-1} + \alpha_{ij} \varepsilon_{i,t-1} \varepsilon_{j,t-1} \quad \forall i, j = 1,\ldots,N \quad (6)$$

where N denotes the number of variables being studied and ω, α and β are parameters (time invariant) of three NxN matrices. Clearly, the VECH model is quite intuitive since it is a simple ARMA process in $\varepsilon_{i,t}\varepsilon_{j,t}$. Although the VECH model is the most general of all the multivariate models and thus provides the greatest flexibility, its complicated structure leads to difficulties in estimation. One of the most important properties for a covariance matrix is positive definiteness. In order for the VECH model to produce a positive definite covariance matrix, many restrictions must be placed on the estimated parameters (Engle et al. 1996).

An attempt to alleviate the problem of positive definiteness was proposed by Engle et al. (1995) in their multivariate BEKK model. The structure of the resulting covariance matrix is as follows:

$$H_t = \Omega + B^T H_{t-1} B + A^T \varepsilon_{t-1} \varepsilon_{t-1}^T A \quad (7)$$

where A, B and Ω are constant matrices of size NxN and the superscripted T denotes the transpose of a square matrix. The rationale behind this approach is that H_t is guaranteed to be positive definite if Ω is positive definite. No further restriction needs to be imposed on the matrices A and B since their use as an outer product formulation produces a quadratic form and therefore is positive definite by construction (Kroner and Ng. 1995). The problem associated with the BEKK representation is the large number of parameters that need to be estimated. Therefore, the use of the BEKK model in a very large financial setting is severely restricted.

As a means of circumventing the problems associated with large-scale estimations, Engle et al. (1990) introduced the Factor ARCH (FARCH) model. This model was developed in the same spirit as the CAPM or the APT pricing models, such that a small number of factors act as forcing variables and drive all of the conditional variances as well as the conditional covariances of asset returns. The individual elements of the covariance matrix under a single factor FARCH model can be found below:

$$h_{ij,t} = \sigma_{ij} + \lambda_i \lambda_j h_{p,t} \qquad (8a)$$
$$h_{pt} = \omega_p + \beta h_{p,t-1} + \alpha \varepsilon_{p,t-1} \qquad (8b)$$

where $i,j = 1,\ldots,N$ and h_{pt} is the conditional variance of the underlying factor (i.e. the market portfolio) which obeys a univariate GARCH(1,1) process. Clearly, the key element under the above FARCH formulation is the selection of the appropriate factor that drives the covariance matrix. In the case of an equity portfolio, the obvious choice would be the market return. In the case of spot foreign exchange returns, the selection of the appropriate factor is not as clear.

Finally, the fourth multivariate GARCH model that we will be discussing here is the constant correlation model (CCORR) of Bollerslev (1990). The general form for the elements of the covariance matrix is given below:

$$h_{ii,t} = \omega_{ii} + \beta_{ii} h_{ii,t-1} + \alpha_{ii} \varepsilon_{i,t-1}^2 \qquad (9a)$$
$$h_{ij,t} = \rho_{ij} \sqrt{h_{ii,t} h_{jj,t}} \qquad (9b)$$

Given the relative simplicity of the above expression, the CCORR model is very easy to estimate due to the small number of parameters involved. Contrary to the previous models discussed and as its name implies, the CCORR model assumes time-invariance for the correlation term and therefore will not be able to capture the heteroskedasticity (if any) in the co-movement of asset prices.

2.2.3. Intradaily Foreign Exchange Volatility

With the advent of advanced data gathering and storage technologies, researchers have taken advantage of the accessibility of intraday spot foreign exchange quotations to further advance our knowledge on the components of the market microstructure. The importance of this data availability lies in the increased statistical significance that arises from large datasets as well as the ability to analyze the behaviour and finer details of the various market participants (Dacorogna et al. 1993).

The foreign currency market is in operation 24 hours a day, seven days a week. In other words, there are no business hour limitations and any market maker can quote a bid/offer price at any point in time during the week. Given this infrastructure, the foreign exchange market can be viewed as the closest proxy to continuous time trading that exists in the world today (Müller et al. 1996(a)). As the research community scrutinized intradaily data on spot foreign exchange quotations, certain characteristics of the price formation process became readily apparent. In Müller et al. (1996(a)), the authors study numerous currencies at various intraday frequencies and find results that are consistent with studies at the daily level. More specifically, the authors confirm the excess kurtosis that has become a trademark of the foreign currency market. Furthermore, their findings indicate that as the sampling frequency increases at the intraday level, the level of kurtosis also increases, which implies a non-convergence of the fourth central moment for the process generating the returns. The authors also measured the level of skewness at various frequencies and found that in all cases, the absolute value of the skewness coefficient to be significantly less than one. Thus implying that the empirical distributions are almost symmetric about the mean. Müller et al. (1990) and Guillaume (1995) also confirmed these results.

In addition to the empirical results mentioned above, strong seasonality patterns were exhibited by most currencies and it was found that this phenomenon could be attributed to periods of low and high trading activity across days of the week. Furthermore, the inclusion of data acquired on the weekends further exacerbated the seasonality patterns (Müller et al. 1996(b)). Initially, excluding the quotations recorded on weekends seemed like the most obvious approach to take given the circumstances. This new time scale is commonly referred to as business time, since it reflects prices recorded during a regular working week. Unfortunately, time series methodologies require equal spacing of the data points and excluding two full days of activity produces a problem in model definition. In order to remedy the problems associated with market inactivity, a new time scale was developed, referred to as θ-time. Essentially, this new time scale contracts physical time during periods of high trading activity and expands it during periods of low trading activity. The rationale behind this new time scale is to remove the seasonal heteroskedasticity associated with changing market activity since each interval of time is constructed to expect the same trading activity. For a more thorough treatment of the θ-time scale see Dacorogna et al. (1993).

Given that the empirical results of intraday data are closely aligned with the documented results at the daily level, it would seem that GARCH is an ideal candidate for modeling the heteroskedastic behaviour of intradaily foreign exchange returns. Unfortunately, of the intraday studies that have attempted to incorporate a standard

GARCH variance structure, none have provided encouraging results. Furthermore, there is no overall consensus in the literature as to which GARCH model is better suited to high frequency foreign exchange returns, or for that matter, even if GARCH is appropriate. The lack of agreement amid the research community is echoed clearly in Müller (1996 (a)). The authors begin by stating that ARCH-type processes are "better suited to capture the tail behaviour in foreign exchange returns, in comparison to simple unconditional volatility models". Although in a subsequent statement, they mention that their analysis of the data leads them to believe that the fourth central moment of the intraday return process is non-convergent – which is in direct conflict with the standard GARCH structure since it relies heavily on the existence of the fourth moment (Diebold 1988).

In order to resolve this apparent dichotomy, we must first understand the differences between volatility defined at the daily level (i.e. where GARCH has been quite successful) to volatility defined at the intradaily level. As stated previously, intradaily foreign exchange returns exhibit seasonal patterns, more so than at the daily level. This has been attributed by numerous studies (Dacorogna et al. 1993 and Guillaume et al. 1995) to the geographic dispersion of the various market agents throughout the world. Therefore, the seasonal heteroskedasticity included in intraday price quotes is more a by-product of the revolution of the earth than one of information flow, one of the main theories used to explain the volatility clustering evident in foreign exchange returns (Tauchen and Pitts 1983). In other words, whereas standard GARCH models were developed to account for the latter, de-seasonalization of the data is a necessary ingredient to a successful study of intraday returns.

In this vein, Guillaume et al. (1995) empirically tested the performance of a univariate GARCH(1,1) process at various intradaily frequencies. Their results indicate that even under a de-seasonalized time scale (i.e. θ-time), the GARCH(1,1) process's predictive power left much to be desired. Furthermore, the authors find that the temporal aggregation properties of GARCH seem to break down at the intradaily level, a result that indicates the presence of various time-horizon components. This finding suggests that the standard GARCH structure is inappropriate for intraday analysis.

As an outgrowth of Guillaume's results, Müller et al. (1996(b)) conducted an exhaustive study of volatility definitions using differing time grids. Their findings suggest that the study of foreign exchange volatility defined over various intraday levels reveals the presence of a heterogeneous market – one that is comprised of market agents with markedly different risk tolerances and time horizons. In order to gain an intuitive understanding of a "heterogeneous market", the authors contrast the intraday foreign exchange trader and a central bank. Intraday price moves go largely unnoticed by the latter, but are considered important events for the former. The author's proceed to test the HARCH (Heterogeneous ARCH) model against the standard GARCH(1,1) and find that HARCH does a better job at describing intraday volatility patterns.

3. Empirical Results

In this section, we will present the empirical findings of our study. Included are the results of the estimated models as well as their performance in the context of Value at Risk. Furthermore, as the empirical evidence pertaining to intraday volatility models in a multivariate framework is sparse, the reader will notice that we have attempted to mimic the methodologies employed in the univariate case in hopes of providing some empirical evidence on foreign exchange volatility in a generalized setting.

3.1. Description of the Study

The purpose of this study is to examine the heteroskedastic nature of foreign exchange returns in a multivariate setting. More specifically, we will investigate the General Dynamic Covariance (GDC) model's (Kroner and Ng 1995) ability to capture the ARCH effects exhibited by foreign exchange returns at the daily as well as at intradaily sampling frequencies. Furthermore, we will be comparing the results of the GDC model with those of the less sophisticated Constant Correlation model (CCORR) of Bollerslev (1990). The two models will then be compared for adequacy on an in-sample basis, with particular attention being paid to the intraday results. We will then proceed to test the performance of these models with respect to their Value at Risk (VaR) forecasts and compare these results to JPMorgan's RiskMetrics™ methodology. The purpose of selecting the CCORR model as a basis for comparison was primarily due to the study by Fallon (1996). Fallon found that the CCORR model produced better estimates of portfolio VaR than models using: 1) a multivariate normal, 2) an IGARCH and 3) an EGARCH model for the state variables. Therefore, this study is aimed at building on the existing evidence.

3.2. The Dataset

The data used in the analysis consists of one year's worth of spot foreign exchange prices for the USD/DEM and the USD/JPY. The data was obtained from Olsen & Associates and was recorded at half-hour intervals starting on January 1, 1996 and running through December 31,1996. All ticks were recorded based on Greenwich Mean Time (to avoid the problems associated with daylight savings time) and included both bid and offer prices. In keeping consistent with the existing literature on high frequency data analysis, we estimated the true[5] price of the exchange rate as follows:

$$x(t_j) = \frac{p_{bid,j} + p_{offer,j}}{2} \qquad (10)$$

where $p_{bid,j}$ and $p_{offer,j}$ are the recorded bid and offer prices at the t_j^{th} tick mark. The return on the spot price was computed as the logarithm of the relative change in the price.

The original data series was recorded in physical time, whereas the time series used in our analysis was actually recorded in business time. In other words, all quoted prices recorded on a weekend were ignored, although prices recorded on holidays were included to avoid the problem of differing holidays in such a global marketplace. We chose not to employ the θ-time scale for two reasons:

- Previous univariate studies of GARCH volatility estimates that employed the θ-time scale did not provide very encouraging results. Therefore, the incremental gain (if any) in using this time transformation does not warrant the added complexity of the procedure for our purposes.

[5] By "true" we refer to the price at which two parties agree to transact.

- From Lundin (1998), the θ-time scale may not be appropriate in a multivariate context since its purpose is to remove selected seasonalities from the data and therefore may eliminate part of the object of measurement.

In our analysis five different series were used, one at the daily level and the others at the intradaily level. They consist of quotes recorded at: 1) half hour, 2) two hour, 3) six hour, 4) twelve hour and 5) daily intervals for both the USD/DEM and USD/JPY rates.

3.3. Preliminary Statistics

Prior to attempting the estimation of the entire covariance matrix, a preliminary analysis of the data was undertaken for the five sampling frequencies in order to gain an intuition on the behaviour of the sample.

Referring to Appendix A, Table A1, we provide sample statistics on the log-relative change of both the USD/DEM and USD/JPY rates. Both of the series were originally demeaned (equivalent to regressing on a constant), since the structure of the mean equation is of little consequence in the subsequent analysis. As stated previously, opinions vary on the actual shape of the return distribution in the currency markets, although there is widespread agreement that foreign exchange returns possess fat-tails. This becomes immediately clear from the results in Table A1. Our results also indicate that the level of skewness of the data, at all sampling frequencies, is substantially less than one in absolute value, which concurs with the previous findings of Westerfield (1977) and more recently of Müller et al. (1996(a)). Therefore, we conclude that our data is almost symmetric with increasing kurtosis as the sampling frequency increases.

One of the main reasons for implementing a GARCH process in the study of foreign exchange returns is to account for the heteroskedasticity inherent in the data. Our next step was to test our sample for evidence of non-constant variance and we follow a similar methodology as Bollerslev (1986) and Hsieh (1989). We employ the popular Ljung-Box statistic to measure the degree of autocorrelation in squared returns. These statistics can be found in Table A2 – Appendix A, for both series as well as the product of the two. The purpose of the latter was used to test the hypothesis of autocorrelation in the covariance term. Given these results, an interesting observation can be made concerning the stationarity of the given sample. For both exchange rates, the hypothesis of homoskedasticity is rejected at the 2.5% level of significance at which point the Ljung-Box statistic tends to decrease quite rapidly. We see similar results for the cross term, but the tests we ran produced much larger statistics indicating to us that the non-stationarity in covariance may also be present. An important point to note however, is that the computed statistics were very low in comparison to previous studies. Hsieh (1989) used 10 years worth of daily data and found the Ljung-Box for up to 50 lags for the USD/DEM to be 215.06. Furthermore, he also computed a value of 206.09 for the USD/JPY with the same amount of data. One potential explanation that may account for the differences in our empirical results is that longer spans of time capture more "regime shifts" in the underlying price process. When considering our results from the higher frequency data, it is clear that these shorter time intervals capture the much shorter-lived movements in the markets. In other words, there appear to be regime shifts on an intraday basis. This result is confirmed in a paper by Müller et al. (1996) where the authors investigate the differences in volatility estimates over various time resolutions. The two key points emphasized in this paper are that the interval size over which volatility is measured is of utmost importance and that volatility measured with high-resolution data contains information not covered by low-resolution estimates.

Given the results of the Ljung-Box test on the covariance term, we then proceeded to analyze the correlation properties of our data in the same way as Lundin et al. (1998). The procedure we used was to take our data and divide it into equal sub-intervals of time. For each of these sub-intervals, we then proceeded to compute correlation estimates between the foreign exchange return pair. The aim of this approach is to study the structure of the correlation coefficient as the time interval being used gets smaller. Our results are displayed graphically in Figure 1 in Appendix A. As is clear from the plots of the correlation coefficients, the estimates appear quite stable when measured over longer intervals of time (see Figure 1(a)). But as the time interval decreases, this stability begins to dissipate (see Figures 1(b)-1(e)) at which point it approaches white noise behaviour in the limit. Lundin et al. refer to this phenomenon as "correlation breakdown". Given this result, the importance of high frequency data becomes readily apparent, since coarsely defined volatility estimates (based on low frequency data) seem to incorporate an averaging effect whereby information contained in the finer time grids is lost.

3.4. The General Dynamic Covariance Model

The underlying motivation behind the development of the GDC model was to describe, in a generalized fashion, the structure of a conditional covariance matrix. As revealed in the previous chapter, a myriad of models have been developed to address the problem of a time varying covariance matrix. Upon further inspection of these model specifications, it is clear that they differ in two very important ways. First, they very different restrictions on the behaviour of the variances and covariances and second, the way they allow the history of the process to affect the variances and covariances (Kroner and Ng 1995). This having been said, the need for a more generalized approach became readily apparent since one particular specification may work well in one circumstance, but not as well in others. The components of the conditional covariance matrix, as specified by the GDC model are provided below:

$$h_{11,t} = \omega_{11} + \beta_{11}^2 h_{11,t-1} + 2\beta_{11}\beta_{21} h_{12,t-1} + \beta_{21}^2 h_{22,t-1} + (\alpha_{11}\varepsilon_{1,t-1} + \alpha_{21}\varepsilon_{2,t-1})^2 \quad (11\,a)$$
$$h_{22,t} = \omega_{22} + \beta_{12}^2 h_{11,t-1} + 2\beta_{12}\beta_{22} h_{12,t-1} + \beta_{22}^2 h_{22,t-1} + (\alpha_{12}\varepsilon_{1,t-1} + \alpha_{22}\varepsilon_{2,t-1})^2 \quad (11\,b)$$
$$h_{12,t} = \sqrt{\theta_{11,t}\theta_{22,t}} \cdot \rho + \sqrt{\theta_{12,t}\theta_{21,t}} \cdot \varphi \quad (11\,c)$$
$$\begin{aligned}\theta_{12,t} &= \omega_{12} + \beta_{11}\beta_{12} h_{11,t-1} + (\beta_{12}\beta_{21} + \beta_{22}\beta_{11}) h_{12,t-1} \\ &\quad + \beta_{21}\beta_{22} h_{22,t-1} + (\alpha_{11}\varepsilon_{1,t-1} + \alpha_{21}\varepsilon_{2,t-1})(\alpha_{12}\varepsilon_{1,t-1} + \alpha_{22}\varepsilon_{2,t-1}) \quad (11\,d)\end{aligned}$$
$$\begin{aligned}\theta_{21,t} &= \omega_{21} + \beta_{11}\beta_{12} h_{11,t-1} + (\beta_{12}\beta_{21} + \beta_{22}\beta_{11}) h_{12,t-1} \\ &\quad + \beta_{21}\beta_{22} h_{22,t-1} + (\alpha_{11}\varepsilon_{1,t-1} + \alpha_{21}\varepsilon_{2,t-1})(\alpha_{12}\varepsilon_{1,t-1} + \alpha_{22}\varepsilon_{2,t-1}) \quad (11\,e)\end{aligned}$$

Clearly, the GDC model possesses a very complicated structure. The reason for this is an outgrowth of the generalized nature of the model. In their paper, the authors provide a list of restrictions to the model that, when imposed in specific combinations, will produce one of the more traditional multivariate GARCH processes. For example, if α_{ij} ,β_{ij} and ϕ are set to zero for i different from j, the resulting conditional covariance matrix is the same as the one specified by the CCORR model of Bollerslev (1990). Similarly, if ρ, α_{ij} and β_{ij} are set to zero for i different from j, then the resulting covariance structure is equivalent to the VECH model (Bollerslev et al. 1988). Cleary, the GDC model does accomplish its primary function, to provide a generalized multivariate GARCH model that acts as a superset to the more traditional parameterizations. Furthermore, the GDC model does not impose one particular structure to the variance and covariance equations, since it inherently captures all of them.

It should also be noted that the authors also consider the problem of positive-definiteness in a covariance estimate. The reader will recall that some of the other multivariate GARCH specifications do not guarantee a resulting estimate that is positive definite. Furthermore, for some of those models, a great many parameters restrictions must be imposed in order to obtain a suitable estimate of the matrix. This requirement, will in most circumstances, introduce estimation problems into the fold. Fortunately, due to the structure of the GDC model, only a size restriction between ρ and ϕ needs to be imposed to guarantee proper behaviour. The interested reader is referred to Kroner and Ng (1995) for a detailed proof.

The complicated structure of the GDC, unfortunately, does not lend itself well to interpretation. Unlike some of the other multivariate GARCH models in use today, the dynamics of the GDC model explicitly incorporate the movements of one factor into those of another. By careful inspection of equations 11(a)-(e), the reader will notice that the volatility of one factor will directly affect all of the others specified in the covariance matrix (albeit a lagged effect - p periods in the general case). This relationship is parameterized through both the covariance term and through the variance equation. In direct contrast, the Constant Correlation model of Bollerslev (1990) makes no such assumption. The CCORR model implicitly assumes that factor volatility is independent across the individual factors, but that the co-movement of the factors themselves is defined by a time-invariant correlation coefficient. This is a key distinction that must be made between these two models, since in some cases volatility in one market may have a direct impact on the volatility of another, a circumstance that is unsupported by the CCORR structure, but embedded in the GDC dynamics.

3.5. Model Estimation Results

In this section, we present the statistical results of the conditional covariance matrix estimation under two different models – the GDC model of Kroner and Ng (1995) and the CCORR model of Bollerslev (1990). The estimation of both models made use of the aforementioned USD/DEM and USD/JPY return series and was based on the full sample for the five different sampling frequencies. The parameter estimates for both of the above models were obtained through the use the R.A.T.S.[6] software package, which implements the Bernt, Hall, Hall and Hausman search algorithm (Bernt et al. 1974) for the maximization of the likelihood function. It should also be noted that we impose a Gaussian error structure for both these models and in both cases, the temporal process to be estimated can be expressed in the following form:

$$r_t = \Sigma_t^{1/2} \cdot \varepsilon_t \qquad (12)$$

where Σ_t is the conditional covariance matrix under either the GDC or the CCORR models, r_t is a 2x1 column vector representing the return series and ε_t is a bivariate white noise process with a mean of zero, unit variance and whose components are independent and identically distributed random variables. The latter random vector is commonly referred to as the forcing variable (Enders 1995) and it is this component that we assume to be Gaussian.

3.5.1. General Dynamic Covariance Model – Results

We refer the reader now to Appendix B – Table B1, which contains the maximum likelihood parameter estimates for the GDC model for the various sampling frequencies.

---

[6] R.A.T.S. version 4.0 was used. A sample program is included in the appendix.

The results indicate that the GDC model does a poor job at describing the data. The t-statistics vary widely for the different sampling frequencies with only the model on the 30-minute data producing seven out of fourteen coefficients as significant at the 5% level. Furthermore, the pattern exhibited by the estimated parameters does not follow the behaviour expected prior to the estimation. As the time interval decreases, we would expect the moving average parameters (i.e. the α_{ij} 's) to tend to zero and the autoregressive parameters (i.e. the β_{ij}'s) to tend to one (Guillaume et al 1995). In most of the cases we investigated, this has not been the observation. Furthermore, the persistence in either of the variances does not seem to follow the growth that was expected with high-resolution data. For the variance of the USD/DEM, we find that the estimated persistence is at a maximum based on 30-minute data, but is at a minimum at two-hour intervals. Similar results were found with the USD/JPY rate. From this, we can confirm Guillaume et al. (1995) results of a break down of the temporal aggregation features of GARCH models at the intradaily level[7].

We then proceeded to analyze the standardized residuals inferred by the model in order to assess whether or not the GDC model adequately removed the heteroskedasticity from the underlying series. The residuals were computed according to the following relationship:

$$\Sigma_t^{-1/2} \cdot r_t = \varepsilon_t \qquad (13)$$

That is, if the model adequately captures the clustering and persistence of variance, the resulting residuals should be described by a bivariate standard normal distribution. The results of the analysis can be found in Table B2 in appendix B.

[7] These results have also been documented with de-seasonlized data (cf. Dacorogna et al. 1997)

From the computed characteristics of the standardized residuals in Table B2, we can see that the residuals are decidedly non-normal. Furthermore, this result is confirmed through the use of the Kolmogornov-Smirnov and χ^2 statistics. In every case, the standardized residuals failed to meet the criteria required by these tests. Again we can conclude that the GDC model does a poor job of describing the dynamics of the USD/DEM and USD/JPY exchange rates.

3.5.2. Constant Correlation Model – Results

In contrast to the results of the GDC model, the results of the CCORR model were more in line with what we expected. For most of the cases investigated, the estimated parameters of the model were highly significant. Furthermore, the size of the parameters were homogeneous for the most part, but the degree of persistence in the variance equations were again greater than one – and hence integrated. This finding again corroborates with Guillaume et al. (1995), such that there is a breakdown of the temporal aggregation process at the intradaily level for GARCH processes.

From an analysis of the standardized residuals, we find similar results to those found under the GDC. The residuals are again decidedly non-normal, a result we expected since no steps were taken to de-seasonalize the data from the outset. Although these results are slightly more encouraging, when coupled with the results of the GDC, we can conclude that GARCH processes, even in the multivariate case, fail to adequately capture the various dynamics that are present at the intraday level.

One final test was performed to test the efficiency of the GDC and CCORR models. It has been well documented in the literature that GARCH processes have very little predictive power on an out-of-sample basis. Bollerslev et al. (1992) and Brailsford et al.

(1996) demonstrate that for a volatility model to produce accurate forecasts, the following regression:

$$r_t^2 = \alpha + \beta \cdot h_t + \varepsilon_t \qquad (14)$$

should produce α very close to zero and β near unity. Furthermore, for h_t to be a good predictor of the squared error, the resulting R^2 should be quite high (i.e. high explanatory power). In a recent paper by Andersen et al. (1998), the authors show that judging a GARCH volatility model based on a low R^2 is incorrect. They claim that the low R^2 produced by GARCH models is a by-product of GARCH itself. The authors show, by way of proof, that the coefficient of determination is bounded above by κ^{-1} where κ is the theoretical kurtosis of the underlying error density. For example, in the case of Gaussian errors, the upper bound on R^2 would be 1/3. Clearly, error densities with greater kurtosis will produce a tighter upper bound.

In order to test whether or not the GDC and CCORR models were good predictors of the expected squared returns, we divided our sample into two parts[8] and re-estimated the models. The regression outlined in (14) was then estimated on the second half of the sample (i.e. the actual squared returns) against the estimates of h_{ij}. For the regression entailing the covariance term, we simply replaced the squared returns with the product of the return series. The results of the regressions are displayed in Appendix B – Tables B7 and B8 for the GDC and CCORR respectively. Please note that the results of the maximum likelihood procedure on the first half of the sample are provided in Tables B5 and B6 respectively.

[8] This was done for all sampling frequencies.

The results of the regression provide some interesting results. For the most part, the regression co-efficients for the GDC model are insignificant. Whereas the intercept terms are very close to zero, the slope co-efficients vary widely and do not approach unity. In contrast, the CCORR regressions prove to be highly significant and the estimated regression parameters are in line with the theory. It should be noted though, the R^2's produced are still quite low, even for the intraday samples – a result that does not concur with the findings of Andersen et al. (1997).

3.6 Comparison of Models

In this section we will be comparing the resulting VaR estimates generated by both the GDC and CCORR models. In addition, we will be comparing these results to the VaR estimates produced by the RiskMetrics™ methodology as outlined by the JPMorgan RiskMetrics™ Technical Document (1995).

The VaR methodology that we employ is based on the parametric Variance-Covariance approach with a delta approximation to the portfolio function. Clearly, since our portfolio consists strictly of spot foreign exchange positions, the return distribution is linear in its payoffs with a delta equal to unity for each of the positions. We will then compare the resulting VaR estimates from the three methods across the five sampling frequencies.

The volatility model defined in the RiskMetrics™ documentation that we will be comparing to the GDC and CCORR models to is:

$$\sigma_{it}^2 = \sqrt{(1-\lambda)\sum_{t=1}^{T}\lambda^{t-1}(r_{1t}-\mu_1)^2} \qquad (15a)$$

$$\sigma_{12t} = \sqrt{(1-\lambda)\sum_{t=1}^{T}\lambda^{t-1}(r_{1t}-\mu_1)(r_{2t}-\mu_2)} \qquad (15b)$$

This model is actually a special case of an IGARCH(1,1) with a zero intercept term. It uses a finite historical window from which to draw (T). Note that this is in direct contrast to other GARCH models, since by definition GARCH has infinite memory.

In order to investigate the relative performance of each of these models, we carry out tests similar to those used by Fallon (1996) and Hendricks (1996). Specifically, we wish to ascertain: 1) the differences in size of the estimates produced by the various methods on a given date, 2) the actual performance of the estimate and 3) for those dates that the actual loss breached the VaR estimate, how large was the shortfall.

In order to accomplish this task, a hypothetical portfolio was constructed from the two spot foreign exchange positions. Then, for a particular sampling frequency, the portfolio was subjected to the actual gains/losses over that period. By repeating this procedure[9] until the end of the sampling period, we effectively construct a series of gains/losses and the associated VaR measures. We then repeated this process nine times by varying the initial portfolio allocation so that the results would not be biased against a particular holding.

In order to investigate how the VaR estimates differed in size, we computed a statistic referred to in Hendricks (1996) as the mean relative bias (MRB). This quantity is defined as being the percentage difference between the VaR estimate for a particular model and the average VaR estimate for that period provided by all three models. Then, the average is taken over all of the periods in the sample. This quantity provides us with an understanding of how, on average the VaR estimates differed from the mean. For example, at the daily level, we find the MRB to be 2.33%, -1.43% and –0.9% for the GDC, CCORR and RiskMetrics™ respectively. These numbers should be interpreted as the size of a VaR estimate relative to the overall mean. At the half-hour level, the MRB is computed to be 0.42%, -0.21% and –0.2% for the GDC, CCORR and RiskMetrics™ model respectively. These results lead us to believe that the RiskMetrics™ model lies in between the other two models with the GDC providing the largest forecasts.

One of the tests we performed was to evaluate if each of the models provided adequate coverage in the event of loss. In Appendix C – Tables C1 and C2, we provide a summary of the percentage of losses exceeding the forecasted risk measure for 99% and 95% VaR estimates respectively. Clearly, the GDC model provides us with the best coverage at the intraday level but does not provide adequate coverage at the daily level. Another interesting result is the performance of the constant correlation model at the intradaily level versus its performance at the daily level. The CCORR model performed better, almost two times better at the daily level, when compared to the intraday results. Finally, the RiskMetrics™ model results are for the most part in between the GDC and CCORR's but failing to provide adequate coverage at all sampling frequencies. In addition, Tables C3 through C6 in Appendix C, contain the z-statistics as well as the confidence intervals on the success rates generated by the three models for a confidence level of 95%. The reader will notice that for the most part, the VaR estimates produced by the three models do not differ significantly (in a statistical sense) from the VaR and that the confidence intervals around the true VaR are quite narrow.

The next statistic we were interested in computing complements the measure of adequacy described above. We computed for each VaR model and for each period where

---

[9] A program was written in C to accomplish this task.

there was a breech in the VaR limit, the amount of the shortfall as a multiplier. Then each multiplier was averaged over the entire testing period. What we found, was that the GDC estimates were closer to the true VaR than any of the other methods. The multiplier that was computed for the GDC for all sampling frequencies ranged from 1.45 to 1.04, whereas the results for the CCORR range from a minimum multiplier of 1.02 to a maximum of 2.04 (i.e. a 100% shortfall). The results for the RiskMetrics™ model were the least encouraging, they ranged from 1.45 to 2.30.

4. Concluding Remarks

In this study, we have attempted to address two issues: Value-at-Risk and stochastic covariance both at the daily and intradaily frequencies in foreign exchange market. We evaluated and compared two competing models that describe heteroskedasticity and have found some interesting results. The generalized model, The General Dynamic Covariance Model of Kroner and Ng (1995) did not perform quite as well as expected. The lack of significance in the estimated parameters leads us to believe that the model is not well suited to the foreign exchange market, at least not when used in conjunction with business time. Furthermore, even with its complicated structure, the model did not adequately capture the inter-temporal and seasonal volatility patterns that are a trademark of intraday foreign exchange quotes.

In contrast, the less sophisticated Constant Correlation Model of Bollerslev (1990), performed quite well in the statistical sense. The significance of the parameter estimates as well as the success the model enjoyed with respect to the traditional efficiency tests, at all of the sampling frequencies, leads us to believe that perhaps the constant correlation assumption is not as restrictive as originally anticipated. What remains to be seen is if either the GDC or the CCORR perform better with a non-Gaussian error structure and possibly under a de-seasonalized time scale.

With respect to Value-at-Risk, we would postulate that the results of the performance tests of the three models are mis-leading. Given that the GDC model outperformed both the CCORR and the RiskMetrics™ model at the intraday level, where model misspecification was greatest, this cannot be misconstrued as a state-variable formulation that describes the covariance dynamics effectively. The suspicions surrounding the GDC model and the resulting large VaR estimates are plentiful and well grounded. In practice, over-estimating VaR will lead to risk-based capital allocation policies that are restrictive and could lead to an inability of the firm to compete in the market. Therefore, overly conservative models cannot be accepted without question and should be closely scrutinized prior to implementation.

From a practitioner's standpoint, VaR has developed into a very effective measure of summarizing firm-wide exposure to market risk. In this paper, we have studied two complex formulations of the covariance matrix as well as one method which in contrast, is much easier to implement. The results of our study do not point to one method that is clearly superior in protecting against adverse market movements. This having been said, the practice of Value at Risk in a banking environment requires a trade-off between model adequacy and ease of implementation – an important trade-off given the constraints inherent in the risk management practice.

References

Andersen, Torben G. and Bollerslev, Tim, 1997, "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts", Working Paper, J.L. Kellogg Graduate School of Management, Northwestern University.

Baillie, Richard T. and Bollerslev, Tim, 1989, "The Message in Daily Foreign Exchange Rates: A Conditional Variance Tale", Journal of Business and Economic Statistics, 7 No. 3, 297-305.

Basle Committee on Banking Supervision, 1997, "Core Principles for Effective Banking Supervision".

Beder, Tanya Styblo, 1995, "VaR: Seductive but Dangerous." Financial Analysts Journal pp. 12-24.

Bernt, E.K., Hall, B.H., Hall, R.E. and Hausman, J.A., 1974, "Estimation and Inference in Nonlinear Structural Models", Annals of Economic and Social Measurement, 4, 653-665.

Bollerslev, Tim, 1986, "Generalized AutoRegressive Conditional Heteroskedasticity". Journal of Econometrics 31 pp.307-327.

_____________ and Wooldridge, J, 1988, "A Capital Asset Pricing Model with Time-Varying Covariances", Journal of Political Economy, 96, 116-131.

____________, 1990, "Modeling the Coherence in Short-Run Nominal Exchange Rates: A Multivariate Generalized ARCH Approach", Review of Economics and Statistics, 72, 498-505.

____________, Chou, Ray Y. and Kroner, Kenneth F.,1992, "ARCH Modeling in Finance – A Review of the Theory and Empirical Evidence", Journal of Econometrics, 52, 5-59.

Brailsford, Timothy J. and Faff, Robert W., 1996, " An Evaluation of Volatility Forecasting Techniques", Journal of Banking and Finance, 20, 419-438.

Burt, J., Kaen, F.R. and Booth, G.G., 1977, "Foreign Market Efficiency Under Flexible Exchange Rates", Journal of Finance, 32, 1325-1338.

Dacorogna, Michel M., Müller, Ulrich A., Nagler, Robert J., Olsen, Richard B. and Pictet, Olivier, 1993, "A Geographical Model for Daily and Weekly Seasonal Volatility in the Foreign Exchange Market", Journal of International Money and Finance, 12, 413-438.

Danielsson, Jon and De Vries, Casper G.,1997, "Extreme Returns, Tail Estimation and Value-at-Risk", Working Paper, London School of Economics and Institute of Economic Studies at University of Iceland.

Diebold, Francis X., 1987, "Empirical Modeling of Exchange Rate Dynamics", Springer-Verlag.

Enders, Walter, 1995, "Applied Econometric Time Series", New York: John Wiley & Sons.

Engle, Robert F., 1982, "AutoRegressive Conditional Heteroskedasticity With Estimates of the Variance of United Kingdom Inflation". Econometrica Volume 50 No. 4.

___________, and Rothschild M., 1990, "Asset Pricing with a Factor ARCH Covariance Structure", Journal of Econometrics, 45,213-238.

___________, and Mezrich, Joseph, 1996, "GARCH for Groups", RISK, pp. 36-40.

Fallon, William, 1996, "Calculating Value-at-Risk", The Wharton Financial Institutions Center. Working Paper 96-49.

Guillaume, D.M., Pictet, O.V. and Dacorogna, M.M., 1995, "On the Intra-Daily Perfromance of GARCH Processes", Working Paper, Olsen & Associates Research Group.

Hendricks, Darryl, 1996, "Evaluation of Value at Risk Models Using Historical Data." Federal Reserve Bank of New York Economic Policy Review pp. 39-70.

Hopper, Gregory P., 1996, "Value at Risk: A New Methodology For Measuring Portfolio Risk." Federal Reserve Bank of Philadelphia Business Review pp. 19-31.

Hsieh, David A., 1988, "The Statistical Properties of Daily Foreign Exchange Rates: 1974-1983", Journal of International Economics, 24, 129-145.

___________, 1989, "Modeling Heteroskedasticity in Daily Foreign Exchange Rates", Journal of Business and Economic Statistics, Volume 7, No. 3, 307-317.

Jones, Mark B. and Schaefer, Stephen M., 1997, "Non-Linear Value-at-Risk", Working Paper.

Jorion, Philippe, 1995, "Predicting Volatility in the Foreign Exchange Market", The Journal of Finance, Volume L, No. 2, 507-528.

___________, 1997, "Value at Risk: The New Benchmark for Controlling Market Risk", Irwin Professional Publishing.

Kroner, Kenneth F. and Ng, Victor K., 1995, "Modeling the Time-Varying Comovement of Asset Returns" Working Paper, International Monetary Fund.

Lamoureux, Christopher G. and Lastrapes, William D., 1990, "Persistence in Variance, Structural Change and the GARCH Model", Journal of Business & Economic Statistics, Volume 8, No. 2, 225-232.

Linsmeier, Thomas J. and Pearson, Neil D., 1996. "Risk Measurement: An Introduction to Value at Risk." Working Paper University of Illinois at Urbana-Champaign.

Lundin, Mark, Dacorogna, Michel M. and Müller, Ulrich, A., 1998, "Correlation of High Frequency Financial Time Series", Working Paper, Olsen & Associates Research Group.

Markowitz, Harry M., 1959, "Portfolio Selection: Efficient Diversification of Investments". New York: John Wiley & Sons.

Müller, Ulrich A., Dacorogna, Michel M. and Pictet, Olivier, 1996(a), "Heavy Tails in High-Frequency Financial Data, Working Paper.

______________, Dacorogna, Michel M., Dave, R.D., Olsen, R.B., Pictet, Olivier and Von Weizsacker, J.E., 1996(b), "Volatility of Different Time Resolutions – Analyzing the Dynamics of Market Components. Working Paper, Olsen & Associates Research Group.

Müller, Ulrich A., Dacorogna, Michel M., Olsen, Richard B., Pictet, Olivier, Schwarz, Mathias and Morgenegg, Claude, 1990, "Statistical Study of Foreign Exchange Rates, Empirical Evidence of a Price Change Scaling Law and Intraday Analysis", Journal of Banking and Finance, 14, 1189-1208.

Phelan, Michael J., 1995, "Probability and Statistics Applied to The Practice of Financial Risk Management: The Case of JP Morgan's RiskMetrics™." The Wharton Financial Institutions Center, Working Paper 95-19.

RiskMetrics™ - Technical Document, 1995, 3rd Edition, Global Research, Morgan Guaranty Trust Company.

Rogalski, R.J. and Vinso, J.D., 1978, "Empirical Properties of Foreign Exchange Rates", Journal of Economic Business Studies, 9, 69-79.

Simons, Katerina, 1996, "Value at Risk – New Approaches to Risk Management." New England Economic Review pp. 3-13.

Tauchen, G. and Pitts, M., 1983, "The Price Variability-Volume Relationship on Speculative Markets", Econometrica, 51, 485-505.

Westerfield, J.M., 1977, An Examination of Foreign Exchange Risk Under Fixed and Floating Regimes", Journal of International Economics, 7, 181-200.

Zangari P. and Mason W., 1997, "An Overview of Value-at-Risk", Global Research, Morgan Guaranty Trust Company.

Appendix A: Preliminary Analysis

Table A1: Preliminary statistics for the USD/DEM and USD/JPY rates.

Sampling Frequency	Statistic	USD/DEM	USD/JPY
Half Hour			
	Standard Deviation	0.00067	0.00081
	Kurtosis	9.74344	7.12924
	Skewness	-0.32027	-0.07459
	Correlation	0.45954	
Two Hour			
	Standard Deviation	0.00124	0.00141
	Kurtosis	6.75986	3.55072
	Skewness	-0.14224	-0.13995
	Correlation	0.49875	
Six Hour			
	Standard Deviation	0.00220	0.00250
	Kurtosis	4.93493	5.25829
	Skewness	-0.47043	-0.14633
	Correlation	0.54547	
Twelve Hour			
	Standard Deviation	0.00296	0.00342
	Kurtosis	1.19176	1.62936
	Skewness	-0.21319	-0.15541
	Correlation	0.56030	
Twenty-Four Hour			
	Standard Deviation	0.00412	0.00460
	Kurtosis	2.35997	0.66922
	Skewness	-0.31052	-0.08678
	Correlation	0.49758	

Table A2: Ljung-Box statistics for autocorrelation in squared returns[a].

Series	Half Hour	Two Hour	Six Hour	Twelve Hour	Twenty-Four Hour
USD/DEM	101.0112	27.3900	18.8345	10.2829	6.4048
USD/JPY	222.7682	29.7201	17.8589	7.7245	8.3255
Product	464.6587	143.5378	29.7532	16.6858	3.4210

[a] The χ^2 statistic with 15 degrees of freedom corresponding to a significance level of: 1% (30.58), 2.5% (27.49), 5% (25.00) and 10% (22.31).

Figure A1: Correlation coefficients calculated using increasingly small sub-intervals[b].

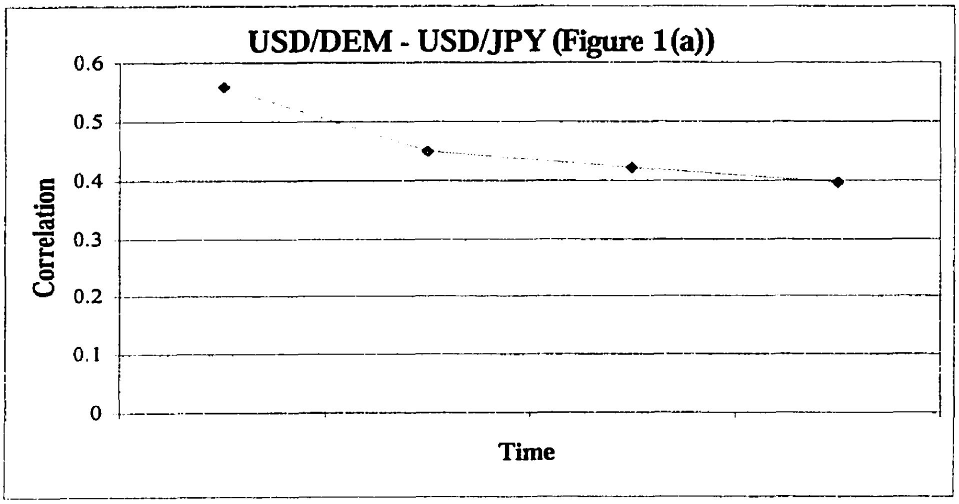

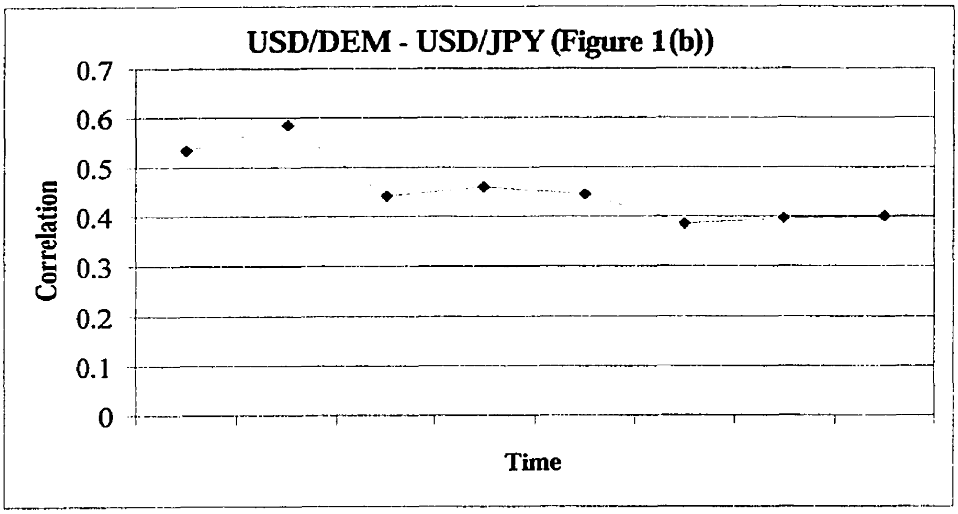

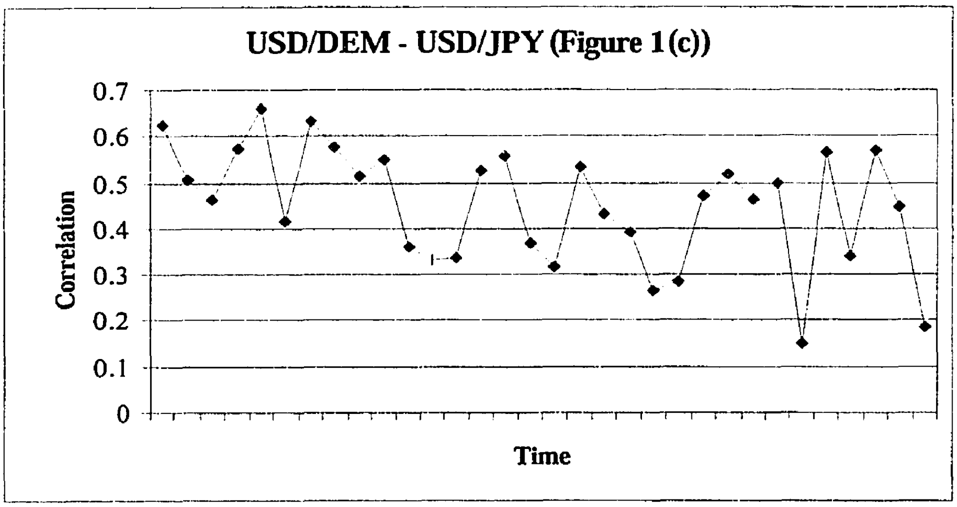

Figure A1: (continued).

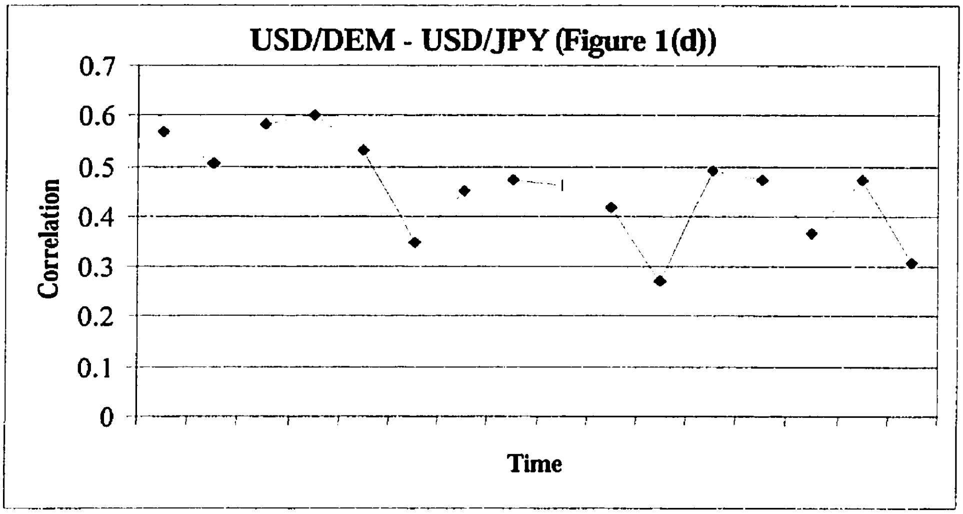

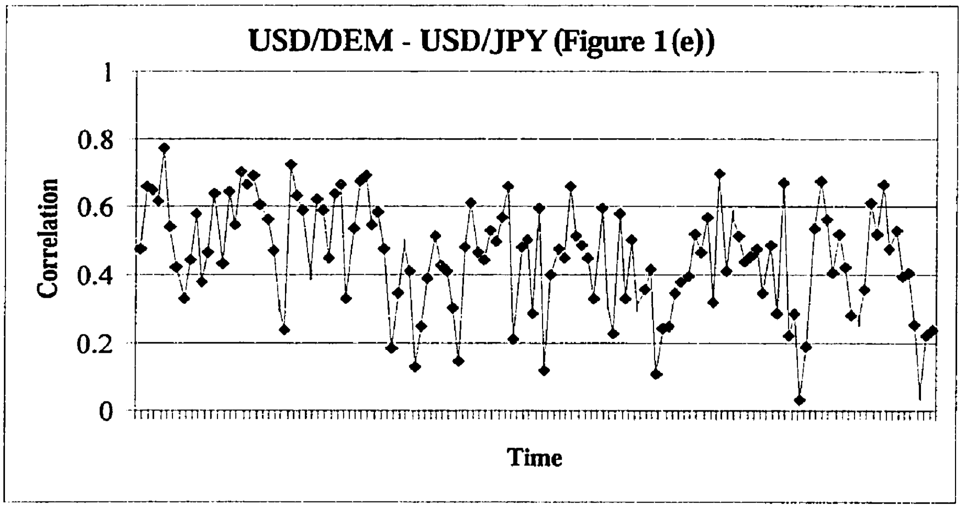

[b] 3132, 1566, 783, 392, 98 data points used for Figures (1a), (1b), (1c), (1d) and (1e) respectively.

Appendix B: Model Estimation Results

Table B1: Maximum Likelihood Estimates – GDC Model (t-statistics[c])

Parameter	Half Hour	Two Hour	Six Hour	Twelve Hour	Twenty-Four Hour
ω_{11}	8.00×10^{-7}	5.00×10^{-7}	3.60×10^{-6}	3.80×10^{-6}	6.45×10^{-6}
	(0.61544)	(0.11254)	(0.15898)	(0.05526)	(0.83262)
ω_{12}	2.00×10^{-6}	1.32×10^{-5}	2.25×10^{-5}	4.24×10^{-5}	6.14×10^{-5}
	(0.74416)	(0.79972)	(0.19350)	(0.12732)	(0.91846)
ω_{21}	3.00×10^{-7}	9.00×10^{-7}	3.00×10^{-6}	5.70×10^{-6}	3.60×10^{-7}
	(0.00)	(0.00)	(0.00)	(0.00)	(0.02851)
ω_{22}	-5.00×10^{-7}	-9.00×10^{-7}	2.40×10^{-6}	8.70×10^{-6}	1.13×10^{-5}
	(-0.40275)	(-0.20243)	(0.11074)	(0.14974)	(1.48646)
α_{11}	9.03×10^{-1}	7.08×10^{-1}	5.92×10^{-1}	6.30×10^{-1}	4.35×10^{-1}
	(12.36944)	(3.30117)	(1.73953)	(1.17558)	(4.26469)
α_{12}	9.65×10^{-2}	7.52×10^{-2}	1.66×10^{-1}	1.67×10^{-1}	1.39×10^{-1}
	(0.97421)	(0.15043)	(0.18347)	(0.10650)	(0.56556)
α_{21}	9.04×10^{-1}	6.75×10^{-1}	6.01×10^{-1}	6.04×10^{-1}	3.39×10^{-1}
	(14.54295)	(4.19154)	(1.78106)	(1.15113)	(3.75695)
α_{22}	1.14×10^{-1}	1.72×10^{-1}	1.34×10^{-1}	2.63×10^{-1}	2.19×10^{-1}
	(3.20309)	(0.45154)	(0.16370)	(0.17485)	(0.99752)
β_{11}	5.89×10^{-1}	6.04×10^{-1}	6.53×10^{-1}	6.92×10^{-1}	5.13×10^{-1}
	(5.84598)	(2.88939)	(1.99989)	(1.03024)	(1.25397)
β_{12}	5.63×10^{-1}	6.13×10^{-1}	6.31×10^{-1}	6.18×10^{-1}	4.51×10^{-1}
	(8.65161)	(4.02776)	(2.81544)	(1.28938)	(0.96921)
β_{21}	7.34×10^{-1}	4.77×10^{-1}	1.64×10^{-1}	2.26×10^{-1}	7.91×10^{-2}
	(3.63007)	(0.95734)	(0.18398)	(0.15557)	(0.09511)
β_{22}	6.01×10^{-1}	5.41×10^{-2}	2.88×10^{-1}	2.56×10^{-1}	3.48×10^{-1}
	(2.66523)	(1.05179)	(0.31860)	(0.18895)	(0.59137)
ρ	-5.06×10^{-2}	-9.06×10^{-4}	-2.50×10^{-1}	-3.46×10^{-1}	-1.24×10^{-1}
	(-0.16611)	(-0.00133)	(-0.18519)	(-0.13203)	(-0.13042)
φ	5.10×10^{-2}	1.33×10^{-1}	6.45×10^{-1}	6.84×10^{-1}	4.60×10^{-1}
	(0.14682)	(0.15742)	(0.37908)	(0.22885)	(0.38951)
Persistence					
USD/DEM	4.844768	3.234666	2.73437	3.016078	1.396037
USD/JPY	1.399235	1.392833	0.93658	0.952027	0.769788

[c] The critical t-statistics for 14 degrees of freedom are 1.761,2.145 and 2.624 at 5%, 2.5% and 1% level of significance respectively.

Table B2: Standardized Residuals – GDC Model

Frequency		Mean	Variance	Kurtosis	Skewness	K-S	χ^2
Half Hour	ε_1	-0.00175	0.10388	7.69167	-0.40815	0.3965	166610.6
	ε_2	0.000434	0.595844	9.35515	-0.08919	0.36032	13692.02
2-Hour	ε_1	-0.001	0.888	7.325	0.118	0.338	1833.956
	ε_2	0.000	2.129	3.534	0.231	0.301	889.386
6-Hour	ε_1	0.001	0.212	2.793	-0.139	0.385	2351.633
	ε_2	-0.001	0.303	3.463	0.295	0.386	2043.118
12-Hour	ε_1	-0.003	0.178	1.413	-0.081	0.396	2075.082
	ε_2	-0.002	0.271	1.417	0.237	0.390	1163.282
24-Hour	ε_1	-0.00042	0.840957	1.447454	0.055436	0.31954	98.96346
	ε_2	-0.004	0.701	0.540	0.162	0.350	110.247

Table B3: Maximum Likelihood Estimates – CCORR Model (t-statistics[d])

Parameter	Half Hour	Two Hour	Six Hour	Twelve Hour	Twenty-Four Hour
ω_{11}	$8.00\text{x}10^{-8}$	$4.6\text{x}10^{-7}$	$1.67\text{x}10^{-6}$	$3.09\text{x}10^{-6}$	$3.68\text{x}10^{-5}$
	(35.634)	(14.182)	(5.130)	(1.705)	(20.737)
ω_{22}	$1.30\text{x}10^{-7}$	$6.0\text{x}10^{-8}$	$7.7\text{x}10^{-7}$	$7.91\text{x}10^{-6}$	$2.77\text{x}10^{-5}$
	(28.540)	(6.185)	(4.569)	(4.165)	(2.230)
α_{11}	$2.17\text{x}10^{-1}$	$1.43\text{x}10^{-1}$	$1.27\text{x}10^{-1}$	$6.56\text{x}10^{-2}$	$3.98\text{x}10^{-2}$
	(36.128)	(11.484)	(8.141)	(2.408)	(4.532)
α_{22}	$1.71\text{x}10^{-1}$	$4.42\text{x}10^{-2}$	$8.34\text{x}10^{-2}$	$1.97\text{x}10^{-1}$	$8.44\text{x}10^{-2}$
	(29.885)	(9.854)	(6.465)	(4.708)	(1.781)
β_{11}	$6.24\text{x}10^{-1}$	$5.54\text{x}10^{-1}$	$5.24\text{x}10^{-1}$	$5.80\text{x}10^{-1}$	-1.0217
	(75.721)	(19.745)	(6.830)	(2.628)	(-133.382)
β_{22}	$6.43\ \text{x}10^{-1}$	$9.25\text{x}10^{-1}$	$7.92\text{x}10^{-1}$	$1.31\text{x}10^{-1}$	$-3.46\text{x}10^{-1}$
	(63.277)	(108.460)	(20.911)	(0.799)	(-0.665)
ρ	$4.51\text{x}10^{-1}$	$4.97\text{x}10^{-1}$	$5.40\text{x}10^{-1}$	$5.63\text{x}10^{-1}$	$5.28\text{x}10^{-1}$
	(91.198)	(46.564)	(26.910)	(17.491)	(8.920)
Persistence					
USD/DEM	0.8419	0.6979	0.6522	0.6458	0.9818
USD/JPY	0.8144	0.9694	0.8761	0.3292	0.2624

[d] The critical t-statistics for 7 degrees of freedom are 1.895,2.365 and 2.998 at 5%, 2.5% and 1% level of significance respectively.

Table B4: Standardized Residuals – CCORR Model

Frequency		Mean	Variance	Kurtosis	Skewness	K-S	χ^2
Half Hour	ε_1	-0.0077	3.2578	9.1416	-0.3841	0.3171	6886.913
	ε_2	0.00155	1.00022	6.03439	0.15698	0.32454	7606.583
2-Hour	ε_1	-0.00088	2.45816	6.82699	-0.07099	0.29739	932.471
	ε_2	-0.00488	1.00013	3.22892	0.17940	0.31664	1164.12
6-Hour	ε_1	0.00720	2.208	2.707	-0.136	0.296	304.556
	ε_2	-0.00337	1.000	3.479	0.293	0.317	374.523
12-Hour	ε_1	-0.00102	2.4328	1.1426	-0.1462	0.342	176.587
	ε_2	0.00005	1.0022	0.9660	0.1567	0.312	111.817
24-Hour	ε_1	-0.00035	0.45224	2.28960	-0.30966	0.36552	231.40533
	ε_2	-0.0005	0.9721	0.4723	0.1468	0.3157	47.5332

Table B5: Maximum Likelihood Estimates (Half of Sample) – GDC Model (t-statistics[e])

Parameter	Half Hour	Two Hour	Six Hour	Twelve Hour	Twenty-Four Hour
ω_{11}	1.54×10^{-6} (0.162830)	1.39×10^{-6} (0.05017)	3.45×10^{-6} (1.23026)	5.64×10^{-6} (0.05345)	6.71×10^{-6} (2.09125)
ω_{12}	-5.00×10^{-7} (-0.075680)	5.50×10^{-7} (0.01507)	3.97×10^{-6} (0.00532)	4.97×10^{-5} (0.10758)	-2.70×10^{-6} (-0.16364)
ω_{21}	2.90×10^{-7} (0.0000002)	9.40×10^{-7} (0.00000)	4.49×10^{-6} (0.00560)	6.30×10^{-6} (0.00000)	2.38×10^{-5} (0.26234)
ω_{22}	-5.60×10^{-7} (-0.073190)	2.11×10^{-6} (0.09500)	1.84×10^{-6} (1.19600)	1.13×10^{-5} (0.13191)	1.32×10^{-5} (1.66479)
α_{11}	5.08×10^{-1} (2.208840)	5.02×10^{-1} (0.61864)	5.38×10^{-1} (4.02088)	6.52×10^{-1} (0.88290)	3.96×10^{-1} (2.93104)
α_{12}	8.31×10^{-2} (0.197900)	9.44×10^{-2} (0.06390)	1.22×10^{-1} (2.68208)	1.94×10^{-1} (0.08939)	2.01×10^{-1} (0.56156)
α_{21}	4.91×10^{-1} (2.595630)	4.98×10^{-1} (1.05604)	4.75×10^{-1} (4.47255)	6.20×10^{-1} (0.83289)	3.38×10^{-1} (2.52234)
α_{22}	1.13×10^{-1} (0.357440)	1.01×10^{-1} (0.13521)	1.23×10^{-1} (3.22933)	3.13×10^{-1} (0.14346)	1.23×10^{-1} (0.35105)
β_{11}	5.79×10^{-1} (1.163620)	6.58×10^{-1} (0.48404)	5.29×10^{-1} (2.41955)	6.90×10^{-1} (0.68415)	4.95×10^{-1} (0.85176)
β_{12}	4.99×10^{-1} (1.148480)	4.95×10^{-1} (0.46030)	3.22×10^{-1} (3.59305)	6.43×10^{-1} (0.87432)	5.82×10^{-1} (1.04521)
β_{21}	8.21×10^{-1} (0.902620)	7.72×10^{-1} (0.30546)	4.03×10^{-1} (0.51831)	2.73×10^{-1} (0.12522)	-6.45×10^{-2} (-0.05392)
β_{22}	6.17×10^{-1} (0.649560)	5.13×10^{-1} (0.23045)	3.08×10^{-1} (0.77984)	1.85×10^{-1} (0.09339)	1.87×10^{-1} (0.21619)
ρ	9.54×10^{-2} (0.071740)	-1.19×10^{-1} (-0.02347)	-2.32×10^{-1} (-0.14010)	-4.96×10^{-1} (-0.11712)	2.82×10^{-1} (0.19535)
φ	9.32×10^{-2} (0.068880)	3.18×10^{-1} (0.05979)	4.03×10^{-1} (0.21336)	8.24×10^{-1} (0.17569)	2.56×10^{-1} (0.15608)

[e] The critical t-statistics for 14 degrees of freedom are 1.761,2.145 and 2.624 at 5%, 2.5% and 1% level of significance respectively.

Table B6: Maximum Likelihood Estimates (Half of Sample) – CCORR Model (t-statistics[f])

Parameter	Half Hour	Two Hour	Six Hour	Twelve Hour	Twenty-Four Hour
ω_{11}	2.51×10^{-7} (7.63915)	7.96×10^{-8} (6.04305)	4.36×10^{-6} (2.68377)	5.53×10^{-6} (0.90018)	3.68×10^{-5} (44.49451)
ω_{22}	2.08×10^{-7} (6.18258)	6.58×10^{-8} (4.21328)	8.70×10^{-7} (2.17147)	5.59×10^{-6} (2.88851)	3.51×10^{-5} (1.96883)
α_{11}	2.03×10^{-1} (20.29387)	3.89×10^{-2} (6.44569)	9.24×10^{-2} (2.09160)	-4.75×10^{-2} (-0.97069)	4.56×10^{-2} (7.36966)
α_{22}	2.16×10^{-1} (21.34070)	4.81×10^{-2} (7.00996)	9.40×10^{-2} (6.55699)	2.48×10^{-1} (3.46212)	6.88×10^{-2} (1.27269)
β_{11}	1.66×10^{-1} (8.99297)	9.06×10^{-1} (66.51821)	-3.53×10^{-2} (-0.12136)	4.33×10^{-1} (0.64865)	-1.05 (-19.18610)
β_{22}	4.71×10^{-1} (36.09647)	9.21×10^{-1} (74.17453)	7.87×10^{-1} (20.61705)	3.51×10^{-1} (2.24737)	-5.42×10^{-1} (-0.79991)
ρ	4.46×10^{-1} (59.13405)	5.14×10^{-1} (31.80875)	5.57×10^{-1} (19.52037)	5.78×10^{-1} (13.33587)	5.01×10^{-1} (6.26292)

[f] The critical t-statistics for 7 degrees of freedom are 1.895,2.365 and 2.998 at 5%, 2.5% and 1% level of significance respectively.

Table B7: Efficiency Regressions – GDC Model (t-statistics)

Sampling Frequency		α	β	R^2
Half Hour	σ^2_{11}	$1.7\text{x}10^{-5}$	0.0964	0.0015
		(2.6171)	(0.4441)	
	σ^2_{22}	$1.5\text{x}10^{-5}$	0.1452	0.0006
		(0.9612)	(0.2920)	
	σ_{12}	$3.72\text{x}10^{-5}$	-0.9395	0.0036
		(0.9730)	(-0.6873)	
Two Hour	σ^2_{11}	$9.2\text{x}10^{-7}$	0.3016	0.0404
		(5.9497)	(8.1449)	
	σ^2_{22}	$1.12\text{x}10^{-6}$	0.6890	0.0106
		(5.8157)	(4.1242)	
	σ_{12}	$2.3\text{x}10^{-7}$	0.4794	0.0341
		($1.1\ \text{x}10^{-7}$)	(0.0643)	
Six Hour	σ^2_{11}	$6.4\text{x}10^{-7}$	0.1868	0.0357
		(0.5642)	(4.4087)	
	σ^2_{22}	$2.73\text{x}10^{-6}$	0.1069	0.0050
		($1.36\ \text{x}10^{-6}$)	(0.0656)	
	σ_{12}	$1.2\text{x}10^{-6}$	0.0663	0.0004
		(0.5119)	(0.4863)	
Twelve Hour	σ^2_{11}	$1.8\text{x}10^{-6}$	0.1319	0.0370
		(0.7668)	(3.1827)	
	σ^2_{22}	$8.9\text{x}10^{-6}$	0.0168	$8.3\text{x}10^{-5}$
		(1.7951)	(0.1475)	
	σ_{12}	$-9.0\text{x}10^{-7}$	0.1868	0.0042
		(-0.1617)	(1.0532)	
Twenty-Four Hour	σ^2_{11}	$1.7\text{x}10^{-5}$	0.0964	0.0015
		(2.6171)	(0.4441)	
	σ^2_{22}	$1.5\ \text{x}10^{-5}$	0.1452	0.0006
		(0.9612)	(0.2920)	
	σ_{12}	$3.7\text{x}10^{-5}$	-0.9395	0.0036
		(0.9730)	(-0.6873)	

Table B8: Efficiency Regressions – CCORR Model (t-statistics)

Sampling Frequency		α	β	R^2
Half Hour	σ^2_{11}	2.0×10^{-7}	1.1700	0.0647
		(16.2614)	(29.5017)	
	σ^2_{22}	1.0×10^{-7}	0.7771	0.0426
		(4.9855)	(23.6537)	
	σ_{12}	1.0×10^{-7}	0.9267	0.0102
		(8.2876)	(11.4082)	
Two Hour	σ^2_{11}	7.7×10^{-7}	1.1264	0.0258
		(6.6135)	(9.1236)	
	σ^2_{22}	1.0×10^{-8}	0.9956	0.0309
		(0.0580)	(10.0127)	
	σ_{12}	1.4×10^{-7}	1.3293	0.0099
		(0.9724)	(5.6175)	
Six Hour	σ^2_{11}	1.73×10^{-6}	1.3561	0.0297
		(2.5899)	(5.6594)	
	σ^2_{22}	8.8×10^{-7}	0.8633	0.0201
		(0.6893)	(4.6339)	
	σ_{12}	1.46×10^{-6}	0.7799	0.0036
		(1.7327)	(1.9579)	
Twelve Hour	σ^2_{11}	3.15×10^{-6}	1.5259	0.0105
		(1.2705)	(2.3515)	
	σ^2_{22}	3.76×10^{-6}	0.6708	0.0185
		(1.3954)	(3.1380)	
	σ_{12}	3.20×10^{-6}	0.6700	0.0021
		(1.3349)	(1.0500)	
Twenty-Four Hour	σ^2_{11}	-5.26×10^{-6}	1.3326	0.0338
		(-0.4300)	(2.1333)	
	σ^2_{22}	1.602×10^{-5}	0.1586	0.0002
		(0.7458)	(0.1619)	
	σ_{12}	-2.125×10^{-5}	3.0839	0.0540
		(-1.7701)	(2.7262)	

<u>**Appendix C**</u>: Comparison of Models

Table C1: Fraction of Outcomes Exceeding VaR Limits (99% Confidence)

	GDC		CCORR		RiskMetrics™	
	% Exceeded	Trials	% Exceeded	Trials	% Exceeded	Trials
Half Hour	0.40%	56592	2.60%	56592	1.80%	56592
Two Hour	0.07%	14157	2.80%	14157	1.90%	14157
Six Hour	0.20%	4716	2.39%	4716	2.14%	4716
Twelve Hour	0.04%	2358	2.54%	2358	2.45%	2358
Daily	1.11%	900	1.00%	900	1.77%	900

Table C2: Fraction of Outcomes Exceeding VaR Limits (95% Confidence)

	GDC		CCORR		RiskMetrics™	
	% Exceeded	Trials	% Exceeded	Trials	% Exceeded	Trials
Half Hour	1.78%	56592	2.89%	56592	2.22%	56592
Two Hour	0.16%	14157	3.35%	14157	3.98%	14157
Six Hour	0.46%	4716	2.65%	4716	3.43%	4716
Twelve Hour	1.01%	2358	2.53%	2358	3.82%	2358
Daily	0.66%	900	2.57%	900	3.74%	900

Table C3: Test for Significance of 99% VaR Success Rates (95% Confidence)

	GDC	CCORR	RiskMetrics™
Half Hour	0.0603	0.1608	0.0804
Two Hour	0.0935	0.1809	0.0905
Six Hour	0.0804	0.1397	0.1146
Twelve Hour	0.0965	0.1548	0.1457
Daily	0.0111	0.0000	0.0774

Table C4: Test for Significance of 95% VaR Success Rates (95% Confidence)

	GDC	CCORR	RiskMetrics™
Half Hour	0.1477	0.0968	0.1276
Two Hour	0.2221	0.0757	0.0468
Six Hour	0.2083	0.1078	0.0720
Twelve Hour	0.1831	0.1133	0.0541
Daily	0.1991	0.1115	0.0578

Table C5: 95% Confidence Interval for Successful Coverage (VaR at 99%)

	GDC		CCORR		RiskMetrics™	
	Lower	Upper	Lower	Upper	Lower	Upper
Half Hour	99.55%	99.65%	97.27%	97.53%	98.09%	98.31%
Two Hour	99.89%	99.97%	96.93%	97.47%	97.88%	98.32%
Six Hour	99.67%	99.93%	97.17%	98.05%	97.45%	98.27%
Twelve Hour	99.88%	100.00%	96.82%	98.10%	96.93%	98.17%
Daily	98.21%	99.57%	98.35%	99.65%	97.37%	99.09%

Table C6: 95% Confidence Interval for Successful Coverage (VaR at 95%)

	GDC		CCORR		RiskMetrics™	
	Lower	Upper	Lower	Upper	Lower	Upper
Half Hour	98.11%	98.33%	96.97%	97.25%	97.66%	97.90%
Two Hour	99.77%	99.91%	96.35%	96.95%	95.70%	96.34%
Six Hour	99.35%	99.73%	96.89%	97.81%	96.05%	97.09%
Twelve Hour	98.59%	99.39%	96.84%	98.10%	95.41%	96.95%
Daily	98.81%	99.87%	96.40%	98.46%	95.02%	97.50%